ICANN75 | AGM – RSS Information Session
Tuesday, September 20, 2022 – 09:00 to 10:00 KUL

UNIDENTIFIED FEMALE: Hello, and welcome to the RSS Information session. Please note that this session is being recorded and is governed by the ICANN Expected Standards of Behavior. During this session, questions or comments submitted in the chat will be read aloud if put in the proper format—and I will note though the format soon.

If you would like to ask a question or make a comment verbally, please raise your hand. When called upon, kindly unmute your microphone and take the floor. Please state your name for the record and speak clearly at a reasonable pace. Mute your microphone when you are done speaking.

This session includes automated real-time transcription. Please note this transcript is not official or authoritative. To view the real-time transcription, click on the Closed Caption button in the Zoom toolbar.

To ensure transparency of participation in ICANN's multistakeholder model, we ask that you sign into Zoom sessions using your full name, for example, a first name and last name or surname. You may be removed from the session if you do not sign in using your full name. With that, I will hand the floor over to Fred Baker. You may begin.

FRED BAKER: Hi there. I'm the chair of the RSSAC. Unfortunately, I'm not able to be on site. But I welcome you to the information session. The session will

cover kind of what it is that we do, what does the root server system do? So with that, let me turn it over to Andrew.

ANDREW MCCONACHIE: Thanks, Fred. My name is Andrew Mcconachie. I work for ICANN Policy Support supporting the RSSAC, the Root Server System Advisory Committee. We used to give this session I guess at every ICANN meeting as part of the how it works, kind of a thread at ICANN meetings, and then COVID came and we stopped. I think we're going to get back into the habit of doing this session quite regularly. But I guess this would be the first time we've given this presentation in about three years.

The format is going to be—so I'll talk for a while about technical stuff related to the root server system, and then Ozan will talk for a while. My colleague here, who is Ozan Sahin, will talk about more policy-related aspects of RSSAC and the RSS. And then after that, we have RSOs, root server operators, in the room and RSSAC members to answer any questions that people have. So if you could please hold your questions until the end, that would be great. And with that, next slide, please.

So this is going to be what we're going to be talking about. First, I'll give an overview of DNS, the Domain Name System, quick explanation of anycast and how that's relevant for the root server system and why it matters, what it does. The root server system today because it has a long history. Then after that, I'll turn it over to my colleague, Ozan Sahin, who will talk about the RSSAC and the RSSAC Caucus and the

evolution of the root server system and its governance model. Next slide, please.

This will be an overview of the DNS. Next slide.

So domain names are just one of the identifiers that are used on the Internet. They're also IP addresses. IP addresses are these identifiers that allow computers to really talk to one another. They're not necessarily intended for humans. IPv4 addresses, of course, are somewhat interpretable by humans, but IPv6 addresses are less interpretable by humans. So we have these handy things called domain names, which are as originally intended, were really supposed to be more for the human side of using the Internet so that people didn't have to remember IP addresses.

Now, all hosts on the Internet do have IP addresses. Not all hosts, of course, have domain names or host names associated with them. This is a bit of a kind of dated view of how domain names and IP addresses work. But it's a good idea to have an understanding of how there's this user interface or human side of naming, and then underneath that there's all these IP addresses, which is what the machines are using to communicate. Next slide, please.

So why do we need DNS? Talking about what I was talking about on the last slide, IP addresses are hard to remember and they change. Originally, there was something called host.txt, which was just kind of a mapping of host names to IP addresses. Eventually, it got too big and we needed something that could be distributed and that could replace host.txt and that could scale. So enter the DNS.

As the Internet has changed over the decades, the kind of the way that IP addresses are used is changed as well. I mean, IP addresses can be shared amongst multiple machines that can be behind a load balancer. So instead of being this kind of one-to-one mapping between domain names and IP addresses, we now have quite possibly one domain name could result in multiple IP addresses, or many machines could have one IP address, or it kind of changes into kind of more of a many-to-many mapping. So that's the more modern way to think about this. Next slide, please.

Now, the Domain Name System is a hierarchy with the root zone at the top. Then beneath that, we have what are called top-level domains, some of the more common ones are listed there, and just kind of collectively referred to as the namespace. There's second, third, you can continue on. You see over to the right, there's this kind of named IP address mapping where we have an example, and then its IP address over to the right.

There are many other mappings besides just what I've been talking about mapping domain names to IP addresses. There are also specific types of DNS records, things for mail servers and reverse lookups and Quad A records for IPv6 and that kind of thing. So the DNS these days does a lot more than just map domain names to IP addresses. It's really integral to the way the Internet functions in a way that wasn't necessarily in the mid 1980s. Next slide, please.

So here's some definitions that are probably good to know in order to really understand what the RSSAC does in the root server system. The first one here is the root service. You can really think about this is what

the root server system does, right? This is the service that it provides. It's really the collective service or collective services provided by all the server instances managed by all the operators.

Then we have the root server system. As opposed to being the service, this is more of the thing. So this is the set of root servers that collectively implements the root servers, and there's a lot of them. We'll get to that later.

Then we have the root zone. This is the DNS zone at the top of the hierarchy that we saw on the previous slide. It has no parent. Yes, it contains all the information necessary to contact the TLD registries underneath it.

Then there's an individual instance, the root server anycast instance. This is just one network location, one location on the network that can respond to DNS queries. One way to think about it is that the root server system is made up of many root server anycast instances. Next slide, please.

Now, these are some of the roles we have. These are again definitions, but these are definitions of roles. There's the root zone administrator. This is the organization responsible for managing the data contained in the root zone, which involves assigning the operators of top-level domains and maintaining their technical administrative details. Right now, this is synonymous with IANA, the Internet Assigned Numbering Authority.

We have the Root Zone Maintainer, and they take that information, and then they format it into a zone file and cryptographically sign it, and then distribute it to the root server operators.

Then we have the root server operators, of which there are 12. They're the ones who are then responsible for taking the zone file for the root zone manager or the root zone maintainer, and then hosting it on their servers. Then taking queries from resolvers around the Internet and responding to them using the data that is in this file. So it's this root zone file that the root server operators host on their servers which then contains all the data that they then respond to the queries with. Next slide, please.

This slide is here to illustrate the difference between the root zone and the root server system. If you remember that slide a few back where I was talking about the hierarchy and the namespace, this is really about the difference between the namespace or the data, the information, and the system that serves that information or responds to queries requesting that information, so to speak. So the root zone is the starting point for how resolvers can find TLD name servers and it's managed by ICANN per community policy. It's compiled and distributed by the root zone maintainer and then given to the root server operators. And then every root server, every root server instance, which then makes up the root server system, has the same zone file and serves the same information.

So the root server system responds with data from this zone file. Currently, there are 26 IP addresses. There are 13 IPv4 and 13 IPv6. There's somewhere over 1500 instances. So, yeah, I'll put some

numbers out there. So there are 12 root server operators. There are 13 what are called root server identifiers, which are lettered from A through M. Each one of those has both an IPv4 and an IPv6 address. Now, each one of those root server identifiers has many copies of itself what we call instances which are distributed around the globe. Right now there's over 1500 of them. That number is kind of steadily increasing all the time. So that's why there's not like an exact number because we could put a number on there, but then by the time we give this presentation, again the number would change. So we just say over 1500 at this point. They all serve the same data, which is from the root zone. Next slide, please.

Okay. This is just kind of an involved slide about—it's meant to place the root server system in the whole structure of the DNS and show the role that the root server system plays and how a typical user would conduct kind of normal things on the web, like go to a webpage, and then the resulting queries, and how they would play out throughout the whole DNS system, including the root server system.

So if we start over on the right, we see we have a user and they would like to go to a website called example.com. We have some information about example.com. We already know its IP address over there, but the user doesn't, so we're going to have to go figure that out by querying the DNS to figure out what the IP address of example.com is. So before they can even really make that HTTP request, before they can contact the server, they need to know its IP address.

So what that user is going to do, that user's computer is going to be configured with something called a recursive name server. That

recursive name server will have a lot of caching built into it. That means it stores information that is previously received for a certain amount of time. But for this example, we're going to assume that the recursive name server doesn't actually have any cache results. So we're going to assume that maybe the recursive name server just got turned on. For whatever reason, its cache is empty and it doesn't know any information, so it's going to have to go out.

So the recursive name server gets this query because the user sends a query for it for www.example.com. The recursive name server is going to have to go out. Because it knows nothing, it doesn't know where .com is. All it knows is that it knows where the root servers are. So we'll go query a root server and say, "Hey, do you know where www.example.com is?" Then the root server will come back with no but I can tell you where .com is. It will tell the recursive name server that information.

So then the name server will then contact the com server, and say, "Hey, do you know where www.example.com is?" and the com name server will say, "No, but I do know where example.com is." Then the root name server will contact the name server for example.com and get the final address for www.example.com, and then tell the user that and the user can then visit the website.

Another piece of this throughout all of these queries and responses—and this is what the keys are signifying on the slide—is that we're assuming in this slide that www.example.com is what's called DNSSEC signed. So in addition to getting information about where these different name servers are, the recursive name servers also getting

cryptographic signatures alongside the records that it needs in order to contact these other name servers and it's doing what's called validating them, where it's performing some cryptographic operations that allow it to determine that this is the correct information, that this information hasn't been tampered with, and then it can be certain that this information is the correct information. So all of that is happening when this user wants to go to www.example.com. I think that's it. Yeah. Next slide, please.

Like in the in the last slide, it's kind of emphasizing the root servers only know what servers need to be asked next, right? Like the root servers in the example previously, they didn't know where www.example.com was. They only knew where .com was. You see they just look at the last label, right? The last label is .com, okay, we'll give you the list of com servers. The last label was .net, we'll give you the net servers, and so on.

Also in the example we just talked about, we made an assumption at the beginning that the root name server had no information stored in it, that the cache was empty. That is rarely the case. The vast majority of queries that users will send to recursive name servers will not need to hit a root server in the end because the information will be cached and the recursive name server will be able to respond very quickly to the user with cached information and not have to go out to either a TLD name server or the root servers. Next slide, please.

So the example I just talked about was really kind of a ideal—I don't want to say old way of thinking about DNS. But, of course, DNS is always changing. So it's still correct to think about that last example,

but there are some modern enhancements and modern refinements to DNS. They're being deployed right now.

We did talk about DNSSEC, the signatures that get returned alongside the data from the authoritative servers. DNSSEC is really all about ensuring the correct data. So the resolver performs validation to ensure that the data that the name server responds with is the right stuff and hasn't been modified in some way.

There are also privacy enhancements being rolled out to DNS. Most of these are between what's called the stub resolver and the recursive resolver, which between the user and the recursive resolver, things like DNS over TLS and DNS over HTTPS. This standards work is still ongoing, although I guess I could say that DoH has received a fair amount of deployment. This is because users don't want people to spy on their queries, they don't want to be surveilled.

There's also query name minimization. Query name minimization, also called QNAME minimization, is a way to only send information that the server needs in order to get to the next server. In the example we had previously, our recursive name server was sending www.example.com all the way up to the root. With query name minimization, the recursive server would only ask for the .com name server. So the root name server would not know that the original query from the user was www.example.com because it would only see a query for .com. I guess I would also say that's in the process of being deployed.

Then finally, anycast. Anycast is a way for a single IP address to have multiple servers behind it, and it's really been key to scaling the root

server system. The root server system really wouldn't be able to scale the way it has scaled without anycast. It both improves latency and resilience in path between recursive resolvers and anycast instances. It also protects against DDoS attacks. We have a section coming up specifically on anycast. So we'll just go on to the next slide. Here we are at the explanation of anycast. Next slide.

Unicast versus anycast. Unicast, the way to think about it is you have one server with one IP address. It's a very basic way to think about it. Unicast you have one server with one IP address. In anycast, you have multiple servers with the same IP address.

The client or the recursive resolver doesn't really know or care in anycast which server it goes to because they're all serving the same stuff. They all have the same data so it doesn't really matter which server they hit. This also has some really important qualities for distributed denial-of-service attacks. Because if an attacker wants to take the service offline … In a unicast world, the attacker only has to bring down one server, right? In an anycast world, the attacker has to bring down a lot of servers. So it's not only good at making sure that the client has a very fast path to a server that serves the information the client wants. But it also ensures that if the system is under attack, it will stay up, it will be resilient. So next slide, please.

Here's some diagrams to kind of explain this a bit better. What you see, this is unicast. You have a source on the left and a destination on the right. And in the middle, we have our very small diagram of the Internet, right? A bunch of routers with links between them in arrows showing which way the traffic flows. In reality, of course, the traffic

would flow both ways, but for these diagrams, it's easier to think about it unidirectionally. So you just see the source goes to the destination like that. Next slide.

Now, in anycast, you see that we have multiple destinations. That original destination in the bottom right that was there in the first slide is still there, only for this source, there's a new destination. And the source can go to any one of these destinations because they're all serving the same information. So it just goes to the one on the bottom left because it's closest. Now, a different source located at a different point in the network would go to a different destination, whichever one is closest. Next slide, please.

Now, here we see anycast under a DDoS attack. One of the anycast nodes—because the attacker has placed themselves on the network in the upper bit, the top right there, all of their traffic just goes to one anycast node. It could be that anycast node just fails because of this. It goes offline or it just can't handle the attack. But that's fine because there are other nodes or other instances available to serve the real traffic. So our source is unhindered by this or source doesn't really care that there's an attack going on because there's another instance available for the source to get their information from. In reality, if it was a really, really good DDoS attack, it may hit many instances. But either way, that attack traffic is going to be distributed across many instances. So in that way, it's much more resilient than unicast. Next slide, please.

I think this is you? Is this still me? Okay, cool. All right, so this one's still me. Next slide, please.

So this shows a bit of the history of the root server system dating all the way back really to the beginning of DNS in 1983. Over time, the number of addresses has increased. Some really important changes over time have been the introduction of anycast and the introduction of IPv6, I would say. You see kind of the development from '83 up to '98, and the number of addresses has been static since '98. There's been 13 root server identifiers since 1998. Next slide, please.

This is a table of the root server identifiers. On the left you see that they're lettered A through M, a.root-servers.net, b.root-servers.net, all the way down to m. Then they all have IPv4 and an IPv6 address. Then on the right, you see the organization that's responsible for managing that root server identifier. So that's the root server operator on the far right. Next slide, please.

This is really meant to show just how distributed. If you go to rootservers.org—this is an interactive map—you can drill down to cities and see where the different instances are. This is kind of on a static slide, it's a little bit less impressive, but I think it's more impressive on the website because you can really see what root server identifiers and how many instances there are in each city. But it's really meant to show that they're just distributed all over the world, every continent, I think, with the exception of Antarctica, although probably someone will correct me and say that there's one in Antarctica. But yeah, they're pretty much everywhere. Next slide, please.

So this diagram shows the entire process from updating information about a TLD to how that then works its way through the whole system

to DNS resolvers. So you just have to read it from the left to the right. TLD operations—imagine that is any TLD operator. They're going to have their name servers that are going to be listed in the root zone. So every TLD has name servers listed in the root zone. And occasionally, those need to be updated, the IP addresses need to be changed or the names of the name servers need to change. Then the TLD operator will contact IANA and have this changed. Then a couple of times a day, IANA—well, no. Sorry, I'm getting ahead of myself. IANA will then record those changes and send them off to the root zone maintainer, and the root zone maintainer a couple of times a day will cryptographically sign a new root zone, basically create a new root zone, and then distribute it to all root server operators who then put it on all of their instances around the world. Then those instances are ready to respond to requests from DNS resolvers anywhere in the world. So that's kind of the whole process of how a TLD operator—how that change of the name server information in the root zone kind of works its way through the system, all the way down to the recursive resolver. Next slide, please.

A little bit more about the root server operators. There are 12 different professional engineering groups that are really focused on reliability and stability to service, accessibility for everyone, all Internet users, and cooperating amongst themselves, and professionalism. It's really a diverse group of organizations. Diverse, technically, kind of organizationally, we have universities, government organizations, private companies, diverse geographically in different countries, and kind of diverse funding models as well because there are different

kinds of organizations with private and public and governmental. Next slide, please.

So this is to show the scope of what root server operators are involved in, really focused on the service, really focused on careful operational evolution of the service. If there's going to be deploying suggested changes, it's going to be done very slowly and deliberately. Really, the focus is on stability, robustness, and ensuring that recursive resolvers around the world can maintain reachability to really multiple instances at all times. The operators are not involved in policymaking. The operators don't modify the data they receive from the root zone maintainer. So it served whatever is given to them. Next slide, please.

There's a lot of coordination of the root server operators. Of course, we have these different kinds of I-star bodies, Internet bodies. There's of course RSSAC at ICANN, there's also the IETF. There's the various kinds of network operating groups and RIRs. RSOs are also active at DNS OARC. So there's a lot of different diverse kind of—I wouldn't necessarily call them bodies—but organizations or conferences. Root server operators can gather and work together, sharing data, and creating a better root server system. Next slide, please.

So this is RSSAC020. I was talking a couple of slides back about how root server operators don't modify the data they receive. They just serve whatever it is from IANA. RSSAC020: Client Side Reliability of Root DNS Data is a document published by the RSSAC, I guess, probably about five or six years ago now, which states this in very explicit terms. So every instance of the root server system serves the same data. That's every anycast instance to the root server system has

the same data and serves the same data. That data originates from the IANA. As we talked about earlier, the DNS is a namespace. It's a hierarchy with a single globally unique root. All clients of the root server system—that means all recursive resolvers—are treated equally. The RSSAC supports to continue deployment of DNSSEC. That's the digital signatures that are served along with the data. Next slide, please.

So over time, certain kinds of myths have cropped up about root server operators or the root server system. This slide is really meant to correct a lot of those myths. I don't hear them as much as I used to. I used to hear these myths more often. I don't know. Maybe these presentations are helping to kind of decry some of these myths. But I'll just walk through them because they at least used to be more common.

Root servers control where Internet traffic goes. That's the first myth here, and that's really not true. It's really routers on the Internet that control where Internet traffic goes. The root servers only have the ability to respond to queries for domain names in return, return with name servers and IP addresses.

Second myth is most DNS queries are handled by root server. Yeah, that's not really true because of caching. So caching, as we talked about in our example, most recursive servers cache answers so there's no need to go out to the root servers all that often.

Third myth is administration and service provisioning of the root zone are the same thing. This is very different. It's IANA that really

administers the root zone. The service provisioning is handled by the root server operators and the root server system.

Another myth is that the idea that a specific root server identity has a specific meaning. So one of them is like more preferred or better than another one, and that's just not true.

Another myth is that there are only 13 root servers. In fact, there's a whole bunch of them. Yeah. They're spread all over the world.

Another myth is that the root server operators conduct operations independently, kind of as if they never talk to one another. That's really not true. They're all in this room together, they talk to one another.

Finally, I guess there was a myth that root server operators only receive the TLD portion of a query. That's maybe the more interesting one because that used to be a myth, although with increased concerns in privacy that is becoming a bit more of a reality when I talk about Q&A minimization, although still I think frequently most of the time the full query makes it all the way up to the root server. And maybe we'll see, as QNAME minimization is deployed more and more, maybe that will change, but right now, I don't know. I don't know precisely what the deployment characteristics are of QNAME minimization, but I'm going to say that it's most queries that make it to the root server system or the entire domain name. Next slide, please.

So this is some guidance for people running networks and how they can hopefully ensure that they have good connectivity to root server and are able to, not just in normal operation but also in cases where

there might be one or two links going down or there might be some kind of partial outage. You still want to make sure that you have good connectivity to root server instance. So that's why you want to have three or four nearby instances, you want to turn on DNSSEC validation in your resolvers. So you're actually checking those digital signatures that are returned with the queries, so you can ensure you're getting the right data.

You may want to participate in the RSSAC Caucus, if you have deep interest in root server system stuff or what the RSSAC or the RSSAC Caucus is doing. You may be interested in hosting an anycast instance. You may be interested in actually having an anycast instance in your network or hosting one yourself. If you are interested in that then you can talk to an RSSAC member after this presentation, or send mail to ask-rssac@icann.org. Next slide, please.

Okay. Now, I think it's over to you, Ozan. Please take it away.

OZAN SAHIN: Thank you, Andrew. Hello, everyone. I'm Ozan Sahin, a member of the ICANN Policy Development Support function, and based in Izmir, Turkey. I'm also supporting the Root Server System Advisory Committee and I'd like to talk about the Root Server System Advisory Committee or RSSAC today and the RSSAC Caucus and the evolution of the root server system. So if you can go to the next slide.

Root Server System Advisory Committee has the role of advising the ICANN community and the ICANN Board on matters relating to the

operation, administration, security, and integrity of the Internet's root server system. This is a narrow scope. Next slide, please.

So this is a committee that produces advice primarily for the ICANN Board but also to other ICANN bodies and some other organizations involved in the overall DNS business. To make a distinction between the RSSAC and root server operators, the root server operators are represented inside the RSSAC, but the RSSAC does not involve itself in the operational matters. Next slide, please.

RSSAC is composed of appointed representatives of the root server operators and there are alternates to these representatives. Also there are some liaisons in the root zone management system in RSSAC. So RSSAC Caucus, on the other hand, is a body of volunteers, subject matter experts. It provides expertise for the RSSAC work. So when RSSAC works on various matters, it is a pool of diverse DNS experts that help RSSAC advance its work. The members of the RSSAC Caucus, there's a Membership Committee that reviews the applications, but eventually, the members are confirmed by the RSSAC based on the Statements of Interest of the applicants. Let's go to the next slide, please.

You see on the screen, the RSSAC leadership, RSSAC chair, Fred Baker, you heard from Fred in the beginning of the call, and RSSAC vice chair, Ken Renard, is with us on the table today. Can we go to the next slide, please?

I mentioned that our liaisons in the RSSAC and who are these. There are incoming liaisons from IANA Functions Operator, Root Zone

Maintainer, the Internet Architecture Board, and Security and Stability Advisory Committee, which is another Advisory Committee in ICANN. RSSAC also has some outgoing liaisons, a liaison to the ICANN Board, liaison to ICANN Nominating Committee or NomCom. And two additional committees that came after the IANA transition in 2016, Customer Standing Committee and Root Zone Evolution Review Committee. Next slide, please.

Currently, RSSAC Caucus has over 100 DNS experts. I guess we have 110 at the moment. Again, they apply via submitting a Statement of Interest and Membership Committee reviews those. The RSSAC Caucus members get credit for their individual work and contributions to RSSAC documents. As I said, these are DNS experts who bring diverse expertise to these publications. This also adds transparency to RSSAC's work because RSSAC itself, as I said, is composed of appointed members from root server organizations. So to join RSSAC, you need to be appointed by one of the 12 root server organizations, but any interest member can apply to join RSSAC Caucus. To apply, you can send your Statement of Interest to rssac-membership@icann.org. Next slide, please.

So these are some of the elements that add transparency to RSSAC and RSOs. To highlight some of them, I already mentioned there's RSSAC Caucus adding transparency to RSSAC's work. RSSAC also has a page, rssac.icann.org, where you can see the minutes of the meetings and the membership, the publications. All these are documented and published. RSSAC, of course, meets at ICANN public meetings. So most recently, RSSAC started conducting its meeting in an open way, so

anyone interested in observing can reach out to support staff and request to the observer at ICANN meetings or join the open meetings like this one at ICANN public meetings.

We also conduct tutorials to engage with the community and inform the community. Also you see operational procedures under which RSSAC function, so it is defined in RSSAC Operational Procedures how RSSAC will conduct its work.

For root server operators, Andrew had mentioned showing a map there's this root-servers.org webpage. You can also see the Root-Ops agendas. They publish sometimes collaborative reports on some of the events. So these are the elements that provide transparency to RSSAC and RSOs. Next slide, please.

I'd like to briefly talk about the root server system evolution. Going to next slide, this is an overview of how this all evolved. Starting with June 2018, RSSAC published two documents, RSSAC037 and 038. We will talk about those quickly. I see we have on the 10 minutes to go to the end of the session.

Then in April 2019, based on these documents, ICANN organization published a concept paper. In January 2020, a Governance Working Group was formed to study these papers and advance the work of evolution of the root server system governance.

Then following and observing this work, RSSAC published in November 2021 two documents to define success criteria for this evolution to be successful. These are called RSSAC058 and 059. Based on the input from these documents in March 2020, the root server

operators, all of them were added to the Governance Working Group. Next slide, please.

An overview of RSSAC037, it defines 11 principles for the operation and evolution of the root server system and proposes an initial governance model for the root server system. It also demonstrates how RSSAC037 model would work through a set of scenarios on designation and removal of operators. Next slide, please.

So these are the 11 principles from RSSAC037 and RSSAC055. For the best interest of time, I'll just go to the next slide. You'll see principles relating to transparency and independence and neutrality. So I will not talk about each of those. Let's go to the next slide.

An overview of 038. These are the recommendations that complement the publication RSSAC037. Basically, RSSAC recommends the ICANN Board to initiate a process to produce a final version of the model that was outlined in RSSAC037. Estimate cost of the root server system in developing the model. And implement the final version of the model based upon the principles of accountability, transparency, and sustainability, and service integrity. Next slide, please.

This is an illustration of the model outlined in RSSAC037. You see the stakeholders are defined as the ICANN community, ITAF/IAB, and the root server operators. You see five functions outlined there. On the bottom-right corner, you see how onboarding and offboarding would work based on the performance metrics and to the degree root server operators meet these performance metrics. Next slide, please.

The Governance Working Group has representatives from the root server operators, Country Code Top-Level Domain Names Supporting Organization or ccNSO. It has representatives from Internet Architecture Board and Registry Stakeholder Group and the Security and Stability Advisory Committee. They're also liaisons from the ICANN Board, the IANA, and Root Zone Maintainer. It is tasked with working out the details of the model outlined in 037.

Also the concept paper that ICANN organization published tasked the Governance Working Group with committing to a timeline, working openly and transparently, and also seeking informed contributions to the working group when necessary, and embracing the principles in RSSAC037, then as resource documents referring to RSSAC037 concept paper and any public comment feedback received and also other relevant RSSAC publications. Next slide, please.

In the timeline, you also saw RSSAC published RSSAC058 and 059 about a year ago, November 2021. RSSAC058 defines the success criteria for this root server system governance structure to be successful. This success criteria form a framework to assess the degree to which any proposed root server system governance structure conforms with the previous RSSAC statements concerning this governance. Then the other document, RSSAC059, it contains the recommendations on how the success criteria should be integrated with the recommendations in RSSAC038. Next slide, please.

These are the five documents that I just mentioned taken as resource documents. There is RSSAC042, A Statement on Root Server Operator Independence. RSSAC047, The RSSAC Advisory on Metrics for the DNS

Root Servers and the Root Server System. RSSAC049, The RSSAC Statement on Joining the Empowered Community, another body that was formed after the IANA transition in 2016, giving ICANN more transparency and accountability.

When we go to the next slide, the other two documents are RSSAC055, The Principles Guiding the Operation of the Public Root Server System, and RSSAC056, RSSAC Advisory on Rogue DNS Root Server Operators.

So these are the five documents. I guess that's the end of my part. Unfortunately, we have only three minutes for the questions. Are there any questions? Shall we do the Zoom first? I see a hand in the in the Zoom Room. Hafiz, please go ahead.

HAFIZ FAROOQ: Hello, everyone. This is Hafiz Farooq, new ICANN Fellow. My question was already partially replied by Fred. My question was about DNS anonymization. Is it an open algorithm available to everyone? If we are using an open algorithm, how are we going to use it for securing the data against the attackers? Are we using any open source algorithm for the anonymization? Or what do we mean by anonymization of the query? Is it done in root server operator level or towards the resolver side? Can you explain more about it? Thank you.

KEN RENARD: This is Ken Renard, one of the root server operators. Are you talking about QNAME minimization?

HAFIZ FAROOQ: Yes.

KEN RENARD: Okay. So that is something done by the resolver. So instead of asking the root the full question, your recursive resolver would only ask the necessary pieces. So that's a modification done only on the recursive server. So that's outside of the root server system. It would be done there.

HAFIZ FAROOQ: So it's actually a minimization.

KEN RENARD: Yes, minimization.

HAFIZ FAROOQ: Okay. Okay, no encryption. Yeah. Okay. All right. Thank you very much.

OZAN SAHIN: Thank you, Hafiz. We have Sávyo next in the queue. Please go ahead, Sávyo.

SÁVYO MORAIS: Hello, that's me. This is Sávyo for the record. I'm also in the Fellowship program. I'm going to make adaption in the question that I always have from my students during the classes that they always ask me why

do most of the root servers are maintained by U.S.-based companies or institutions? So my actual question is why do we still have most of the DNS root servers managed by U.S. institutions? Why it's not still more global, more distributed globally? What's the big problem on changing it or adding more servers from different places? Thank you.

LARS-JOHAN LIMAN: Lars Liman here from Netnod, also one of the root name server operators who's been around also for quite some time.

Most of the root server operators, eight of the 12 operators go back so far in time that back then the Internet was an American thing. So there was precious little Internet operation outside the U.S. Since then, a number of operators have been added. So we are now 12 operators and three of them are based outside the U.S. But since 1998, there has been no process for changing operators. It's not that there is a lack of willing people. It's just there is no process telling how to do that.

There used to be a process but it was a one-man process with the old head of the IANA, Jon Postel. He was in charge of the entire root server system. But he died very suddenly before he managed to transfer this responsibility to something new. He was part of creating ICANN, and ICANN was still just being born when he died. So since then, there has been no process and no one has dared to look at this until the root server operators within RSSAC back in 2015 said this is a process that we probably need to have. So we need to look at the problem and design a new structure.

That work is ongoing now in the Governance Working Group. So one of the outcomes of the Governance Working Group is likely to be a process for adding or deleting root server operators, and it may well be the case that we have too many. We don't really know. But there will also be an assessment function to look at the strategy and architecture of the root server system to look into the distant future and see how we want it to work in the future.

So it's for historic reasons. In more recent times, because we don't have a process to change, there's no way for Netnod to stop doing this if we wanted to. So this is all based in history. We are doing work to put the pieces in place to make changes in the future. But right now, we don't have any. Working on it.

SÁVYO MORAIS: Thank you very much. Well, the history, I think, it's clear. It's an old thing, but thank you for the clarification and the process of the policy.

KEN RENARD: The other thing is that this is a technical infrastructure. Regardless of where the operator is, we have instances around the globe. So the service is available locally pretty much wherever you are.

SÁVYO MORAIS: Yes. But if something happened in U.S., maybe this is the point.

OZAN SAHIN: Thank you. I see we're out of time. To respect the next group that will be using the room, let's do one final question. I see in the queue, the next is Chokri, and then close the queue. I see we may need to keep the presentation a bit shorter next time. Apologies for this short time for questions. Chokri.

CHOKRI BEN ROMDHANE: Okay. Thank you. I'm Chokri from Tunisia. My question is about the alternate DNS that are using some new technologies such as blockchain and other technology in order to provide the DNS service. Do you think that such technology is providing some threat or risk to the root server, or in contrast, it provides some cash or limit traffic to the head root server to do its job in a good way? Thank you.

KEN RENARD: Are you talking about alternate namespaces? That doesn't provide a direct threat to the root server system. We have our namespace, it really is a threat to the Internet user when there's multiple namespaces. What are they going to get based on which namespace they're using? So that's why we believe the IANA is the true source of the root zone. That's the one namespace that we will use for our root service.

CHOKRI BEN ROMDHANE: But this technology, they always considered that the root servers are the main feature or the main segment of the Internet. But they provide

service in other layer. So do you think that such technology can make some risk or threaten to the server functionality?

KEN RENARD: It's really a matter of adoption of those things. Again, one of the purposes of ICANN is to try and keep it globally unique namespace. So it's sort of above our purview here as a root server operator in RSSAC, but it is something that we're following and what I'm looking for the stability of the Internet thing.

OZAN SAHIN: Okay. Thank you for joining us today for this presentation. We're adjourned. Have a good day.

[END OF TRANSCRIPTION]